

CLAIMS

1. A method of queuing packets for processing, the method comprising the steps of:

- 5 a. allocating each received packet to at least one arrival queue;
 - b. placing each packet in the allocated queue if said queue is not full, otherwise dropping said packet;
 - c. scheduling packets from the arrival queue to at least one transfer queue;
 - 10 d. responsive to transfer of a packet to a transfer queue, generating an interrupt;
 - e. responsive to receipt of an interrupt, allocating the packet to one of a plurality of processor queues;
 - 15 f. placing the packet in the allocated processor queue if said queue is not full, otherwise dropping said packet; and
 - g. scheduling packets from the processor queues for processing.
- 20 2. A method according to claim 1 wherein packets are received at an input to a plurality of devices.
3. A method according to claim 1 or claim 2 wherein at least one device has a plurality of arrival queues
4. A method according to claim 3 wherein each arrival
- 25 queue is associated with a traffic class, each packet being allocated to the at least one queue in accordance with the traffic class of each packet.

5. A method according to claim 4 wherein the traffic class is priority information embedded in the each packet.

6. A method according to claim 1 wherein at least one device comprises a plurality of transfer queues.

5 7. A method according to claim 1 wherein the number of transfer queues for each device is less than the number of arrival queues for each device.

8. A method according to claim 1 wherein the scheduling of packets from the arrival queue to the transfer queue is
10 dependent upon one or more of: the traffic profile; the quality of service requirement; or the characteristics of the transfer queues.

9. A method according to claim 1 wherein the transfer queue comprises a device level transfer queue and a
15 processor level transfer queue, wherein the device level transfer queue receives packets from the arrival queue, and the processor level transfer queue receives packets from the device level transfer queue.

10. A method according to claim 9 wherein packets are
20 transferred to the processor level transfer queue from the device level transfer queue whenever there is space in the processor level transfer queue.

11. A method according to claim 10 wherein packets are never dropped from the transfer queue.

25 12. A method according to claim 1 wherein the processor queues are associated with different priorities.

13. A method according to claim 12 wherein the highest priority queue has the lowest drop probability and the lowest latency.

14. A method according to claim 1 wherein responsive to receipt of an interrupt from a device, a packet is removed from the transfer queue of the device and classified.

5 15. A method according to claim 14 wherein the classification is based on a determination of priority.

16. A method according to claim 14 wherein the packet is allocated to a processor queue in accordance with its classification.

10 17. A method according to claim 14 wherein the packet is placed in the allocated processor queue if said queue is not full, otherwise the packet is dropped.

15 18. A system including a processor and at least one device, in which system: packets for processing by the processor are received at an input of the at least one device, wherein the at least one device includes:

a. allocating means for allocating each received packet to at least one arrival queue of the device;

20 b. placement means for placing each packet in the allocated queue if said queue is not full, otherwise dropping said packet;

c. scheduling means for scheduling packets from the device arrival queue to at least one transfer queue; and

25 d. interrupt means, responsive to transfer of a packet to a transfer queue, for generating an interrupt from the device to a processor;

and wherein the processor includes:

e. allocation means, responsive to receipt of an interrupt, for allocating the packet to one of a plurality of processor queues;

5 f. placement means for placing the packet in the allocated processor queue if said queue is not full, otherwise dropping said packet; and

g. scheduling means for scheduling packets from the processor queues for processing.

10 19. A system according to claim 18 including a plurality of devices adapted to receive packets for processing by the processor at inputs thereof.

20. A system according to claim 18 or claim 19 in which at least one device is adapted to provide a plurality of arrival queues

15 21. A system according to claim 20 in which each arrival queue is associated with a traffic class, each packet being allocated to the at least one queue by the allocation means in accordance with the traffic class of each packet.

20 22. A system according to claim 21 in which the traffic class is priority information embedded in the each packet.

23. A system according to claim 18 in which at least one device includes a plurality of transfer queues.

25 24. A system according to claim 18 wherein the number of transfer queues for each device is less than the number of arrival queues for each device.

25. A system according to claim 18 in which the scheduling means is responsive to one or more of: the traffic profile; the quality of service requirement; or the characteristics of the transfer queues.

26. A system according to claim 18 in which the transfer queue comprises a device level transfer queue and a processor level transfer queue, the device level transfer queue being adapted to receive packets from the arrival queue, and the processor level transfer queue being adapted to receive packets from the device level transfer queue.

27. A system according to claim 26 being adapted such that packets are transferred to the processor level transfer queue from the device level transfer queue whenever there is space in the processor level transfer queue.

28. A system according to claim 27 further adapted such that packets are never dropped from the transfer queue.

29. A system according to claim 18 in which the processor queues are adapted to be associated with different priorities.

30. A system according to claim 29 in which the system is adapted such that the highest priority queue has the lowest drop probability and the lowest latency.

31. A system according to claim 18 wherein the processor includes transfer means adapted, responsive to receipt of an interrupt from a device, to remove a packet from the transfer queue of the device, and provide such to a classification means for classification.

32. A system according to claim 31 in which the classification is adapted to be based on a determination of priority.

33. A system according to claim 31 further including means to allocate the packet to a processor queue in accordance with its classification.

34. A system according to claim 31 where placement means are adapted such that the packet is placed in the allocated processor queue if said queue is not full, and otherwise the packet is dropped.

5 35. A device adapted for queuing packets to be processed, the device including:

a. allocating means for allocating a received packet to at least one arrival queue;

10 b. placement means for placing each packet in the allocated queue if said queue is not full, otherwise dropping said packet;

c. scheduling means for scheduling packets from the arrival queue to at least one transfer queue; and

15 d. interrupt means, responsive to transfer of a packet to a transfer queue, for generating an interrupt;

e. allocation means, responsive to receipt of an interrupt, for allocating the packet to one of a plurality of processor queues;

20 f. placement means for placing the packet in the allocated processor queue if said queue is not full, otherwise dropping said packet; and

g. scheduling means for scheduling packets from the processor queues for processing.

25 36. A device according to claim 35 including a plurality of arrival queues

37. A device according to claim 36 in which each arrival queue is associated with a traffic class, each packet being allocated to the at least one queue by the allocation means in accordance with the traffic class of each packet.

38. A device according to claim 35 including a plurality of transfer queues.

39. A device according to claim 35 in which the transfer queue comprises a device level transfer queue and a
5 processor level transfer queue, the device level transfer queue being adapted to receive packets from the arrival queue, and the processor level transfer queue being adapted to receive packets from the device level transfer queue.

40. A device according to claim 39 being adapted such that
10 packets are transferred to the processor level transfer queue from the device level transfer queue whenever there is space in the processor level transfer queue.

41. A device according to claim 40 further adapted such that packets are never dropped from the transfer queue.

15 42. A device according to claim 35 in which the processor queues are adapted to be associated with different priorities.

43. A device according to claim 35 further including transfer means adapted, responsive to receipt of an
20 interrupt, to remove a packet from the transfer queue of, and provide such to a classification means for classification.

44. A device according to claim 35 further including means
25 to allocate the packet to a processor queue in accordance with its classification.

45. A device according to claim 44 in which the placement means are adapted such that the packet is placed in the allocated processor queue if said queue is not full, and otherwise the packet is dropped.